ERDC MSRC PET Technical Report No. 01-05

Performance Comparison of SGI Origin 2800 and SGI Origin 3800 on Application Codes

by

Jeff Hensley
Daniel Duffy
Mark Fahey
Tom Oppe
William Ward
Robert Alter

20 April 2001

Work funded by the Department of Defense
High Performance Computing Modernization Program
U.S. Army Engineer Research and Development Center
Major Shared Resource Center through

Programming Environment and Training

Supported by Contract Number: DAHC94-96-C0002
Computer Sciences Corporation

Performance Comparison of SGI Origin 2800 and SGI Origin 3800 on Application Codes

Jeff Hensley*, Daniel Duffy*, Mark Fahey*,
Tom Oppe†, William Ward*, and Robert Alter‡
Computational Migration Group
U.S. Army Engineer Research and Development Center
Major Shared Resource Center
Vicksburg, MS

Abstract

Benchmark tests using a wide variety of application codes show a consistent improved performance on SGI's Origin 3800 (O3K) in comparison to the SGI Origin 2800 (O2K). Most codes tested ran slightly more than twice as fast on the O3K as on the O2K.

1 Introduction

Benchmarks are widely used in industry, government, and academia to analyze the performance of computer hardware and software. Benchmarking results can be used to target new acquisitions of hardware, to identify proper platforms for specific types of applications, and to help allocate limited resources.

Benchmarks can be useful in quantifying the improvement in system performance after a major hardware upgrade. Such benchmarks not only provide a measure of the increased performance but also provide valuable clues to the effects of future hardware upgrades.

In this spirit, a series of benchmark tests were performed on two SGI systems. The first computer is an SGI Origin 2000 series machine consisting of 128 processors each with a chip speed of 195 MHz. This machine, which is located at the Aeronautical Systems Center (ASC) Major Shared Resource Center (MSRC), will be referred to as the Origin 2000 (O2K)[1]. The Engineer Research and Development Center (ERDC) MSRC has recently installed a 512 processor SGI Origin 3800[2]. This machine is one of the first three such machines in the world. This machine, referred to as the Origin 3000 (O3K), has a faster chip speed of 400 MHz and a greatly improved interprocessor bandwidth with reduced latency[3].

2 Application codes tested

The set of application codes studied in this paper are a subset of codes used in the TI01 Department of Defense (DoD) Benchmark[4]. Based on usage throughout the DoD, these codes represent a wide cross-section of the current computational research being performed for the DoD. With the

*Computer Sciences Corporation
†Data Management Consultants
‡Alpha Data Corporation

assistance of users, representative input decks were developed, tested, and run on both the O2K and the O3K. A list of the codes, the input decks and the number of processors used for each test is provided in Table 1.

The following provides a short description of each of the codes as well as information about the input decks.

2.1 CHARGE

CHARGE is a finite volume electromagnetics solver for calculating the radar cross section of complex scatterers. It is a cell-centered scheme utilizing third-order Van-Leer splitting in space and second-order Runge-Kutta integration for time advancement.

The input deck used for these tests specifies the radar cross section of a sphere. This problem scales nicely on both the O2K and the O3K. Results from other studies show nice scaling on other architectures as well[4].

2.2 COBALT-60

COBALT-60 is a parallel Computational Fluid Dynamics (CFD) code that solves the compressible Euler and Navier-Stokes equations over unstructured grids. COBALT-60 is often used to model fluid flow and turbulence around objects that are traveling through the fluid (*e.g.*, a plane or missile flying through the air).

The three input decks represent problems of varying sizes. The first input, "missile," models the flight of a missile moving at Mach 2.5 and uses a total of approximately 725,000 cells. The second input, "wingflap," models a hinged airfoil and requires nearly 3,000,000 cells. The third input, "trapwing," uses a total of about 7,000,000 cells.

2.3 CTH

CTH explores the effects of strong shock waves on a variety of materials using many different models. CTH has models for a variety of materials: multi-phase, elastic-viscoplastic, porous, and explosive. CTH is one of the most widely used applications in DoD research, including the study of blast waves for the analysis of explosive effects on structures.

2.4 FEMD

FEMD is a parallel wave pseudopotential finite element molecular dynamics code. It can be used to analyze and model materials at the molecular and atomic level. For example, FEMD can be used to study a ceramic/metal interface.

The input deck used for this study describes a calculation for 32 atoms with a total of 116 electrons occupying 64 states.

2.5 FEMWATER123

FW123 is a parallel version of the FEMWATER code, a 3-D finite element computer model for simulating surface and groundwater flow through variable media. The code combines the ability to have a 3-D ground volume of varying materials through which groundwater flows coupled to a 2-D surface water model and 1-D canal structures.

The input deck used in this study models the natural environment in southern Florida, with a complex ground water system coupled to one-dimensional canals and two-dimensional surface

Table 1: Benchmark results.

CODE	Input Data	NP	O2K Time (sec)	O3K Time (sec)	Ratio O2K/O3K	
GAMESS	cycl	16	1632	769	2.12	
GAMESS	cycl	32	895	427	2.11	
GAMESS	cycl	64	516	260	1.98	2.06
GAMESS	hedm	32	9517	4371	2.18	
GAMESS	hedm	64	4997	2356	2.12	
GAMESS	hedm	96	3492	1633	2.14	
GAMESS	hedm	128	2804	1270	2.21	2.17
CTH	arm.t1	64	4177	2034	2.05	
CTH	arm.t1	80	3488	1718	2.03	
CTH	arm.t1	96	3063	1498	2.04	
CTH	arm.t1	128	2745	1248	2.20	2.09
FEMWATER	fred	8	3272	1516	2.16	
FEMWATER	fred	12	3473	1528	2.27	
FEMWATER	fred	16	3013	1418	2.12	
FEMWATER	fred	32	2933	1261	2.33	
FEMWATER	fred	48	2605	1273	2.05	2.17
ICEPIC	dat.128	32	7261	5080	1.43	
ICEPIC	dat.128	64	4653	2154	2.16	
ICEPIC	dat.128	128	2380	916	2.60	2.00
ICEPIC	mitl.100	32	1129	614	1.84	1.84
CHARGE	500.25	8	6352	2918	2.18	
CHARGE	500.25	16	3246	1476	2.20	
CHARGE	500.25	32	1659	817	2.03	
CHARGE	500.25	64	907	490	1.85	
CHARGE	500.25	128	703	330	2.13	2.07
COBALT	missile	64	1009	886	1.14	
COBALT	missile	128	686	536	1.28	1.22
COBALT	wingflap	16	8334*	3664	2.27	
COBALT	wingflap	32	4420*	1733	2.55	
COBALT	wingflap	64	2882*	883	3.26	
COBALT	wingflap	128	1004	533	1.88	2.46
COBALT	trapwing	32	7481*	2859	2.62	
COBALT	trapwing	64	3166*	1840	1.72	2.11
LESLIE3D	128x128x128	8	5711	2747	2.08	
LESLIE3D	128x128x128	16	2926	1406	2.08	
LESLIE3D	128x128x128	32	1612	729	2.21	2.13
FEMD	inp	16	10571	4526	2.34	
FEMD	inp	32	5321	2239	2.38	2.12
NLOM	na825	28	14495	6300	2.30	
NLOM	na825	56	7963	2933	2.71	
NLOM	na825	112	3807	1602	2.38	2.46

water in Dade county. A quick look at the results in Table 1 indicate that this problem does not scale well. At larger numbers of processors the cost of inter-processor communication outweighs any gain from reducing the computation time.

2.6 GAMESS

GAMESS (General Atomic and Molecular Electronic Structure System) is a program for computational quantum chemistry. Material properties can be calculated from first principles. The number, type, and position of atoms in a molecule are specified as initial conditions and GAMESS can compute molecular energies and a variety of wave functions and other quantities.

2.7 ICEPIC

ICEPIC (Improved Concurrent Electromagnetic Particle-In-Cell) is a parallel, 3-D particle-in-cell simulation tool for relativistic problems involving collisionless or low-collisionality plasmas in complex geometries. One use of ICEPIC is to model high power microwave sytems and directed energy devices.

2.8 LESLIE3D

LESLIE3D is a parallel code that performs CFD calculations over structured grids. It performs large eddy simulations and is used to model combustion.

LESLIE3D has been used to analyze reacting flows such as combustion in a gas turbine or combustion instability in ramjet engines. However, the problem solved in the benchmark test is simpler: modeling the flow of a temporal mixing layer formed between two parallel plates that are moving in opposite directions.

2.9 NLOM

NLOM implements the Navy Layered Ocean Model. It performs global and basin-scale ocean modeling and prediction. The benchmark test does an analysis for 3.05 model days on a 1/64 degree 5-layer Atlantic Sub-tropical Gyre region. The problem does make significant resource demands on the computing system: about 12 Gigabytes of memory and 20 Gigabytes of scratch disk space is required.

3 Results

The benchmark test results are summarized in Table 1. All tests, with the exception of some of the COBALT input decks on the O2K, were performed during dedicated time on the machines. Entries marked with an asterisk indicate non-dedicated runs; *i.e.*, runs performed during normal use of the machine. The first column in the table is the name of the code. The second column is a name given to the input deck used. The third column refers to the number of processors used. The fourth and fifth columns give the time in seconds for the test run on the O2K and O3K, respectively. The sixth column gives the ratio of the time on the O2K to time on the O3K. The numbers appearing to the right of the sixth column are a measure of the speedup for each benchmark test. For a given input deck, this is computed by the following formula:

$$\text{speedup measure} = \frac{\sum_{np} T_{O2K} \cdot np}{\sum_{np} T_{O3K} \cdot np}$$

where the sum is taken over each configuration size of processors used. Here, T_{O2K} and T_{O3K} refer to the times on the O2K and O3K, respectively, for a particular test; np is the number of processors used for that given test. The total computed speedup measure for all benchmark tests is 2.15.

On most of the test problems, the O3K showed a fairly consistent level of improvement, running the problems a little more than twice as fast as the O2K. The O3K has a faster clock rate and the expected speedup of calculations based on the ratio of the clock rates is $400/195 \approx 2.05$. But all of these codes require interprocessor communication; that is, there is a certain (nontrivial) amount of time during which the code is not engaged in parallel work. However, what seems to be readily obvious from the results is that the increased bandwidth of the O3K does make a definite impact on the performance. Without the improved bandwidth, the time spent in communication on the two machines would have been roughly equal and the overall speedup of the O3K would have been less than twice when compared to the O2K.

Figure 1 gives a graphical comparison of the O3K versus the O2K for four of the benchmark codes. Each of the four tests uses MPI for message passing. Not only is it evident that the O3K is markedly faster, but the graphs also indicate that each of these codes is scalable for the problems. Also, note that the O3K and O2K scale in the same fashion for these tests.

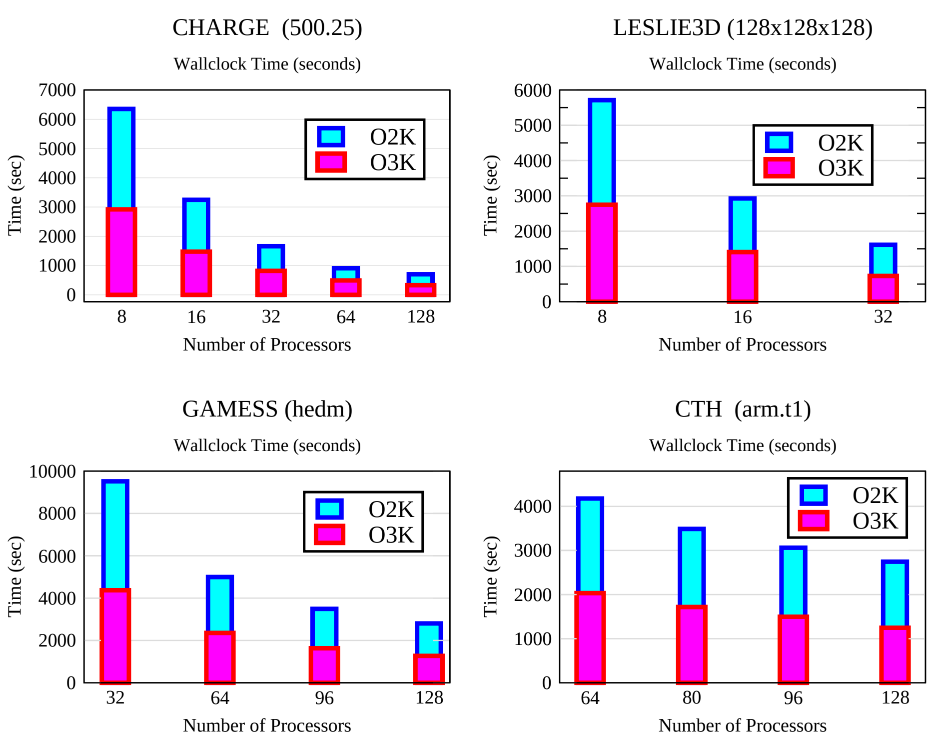

Figure 1: Timing comparisons for four benchmark codes.

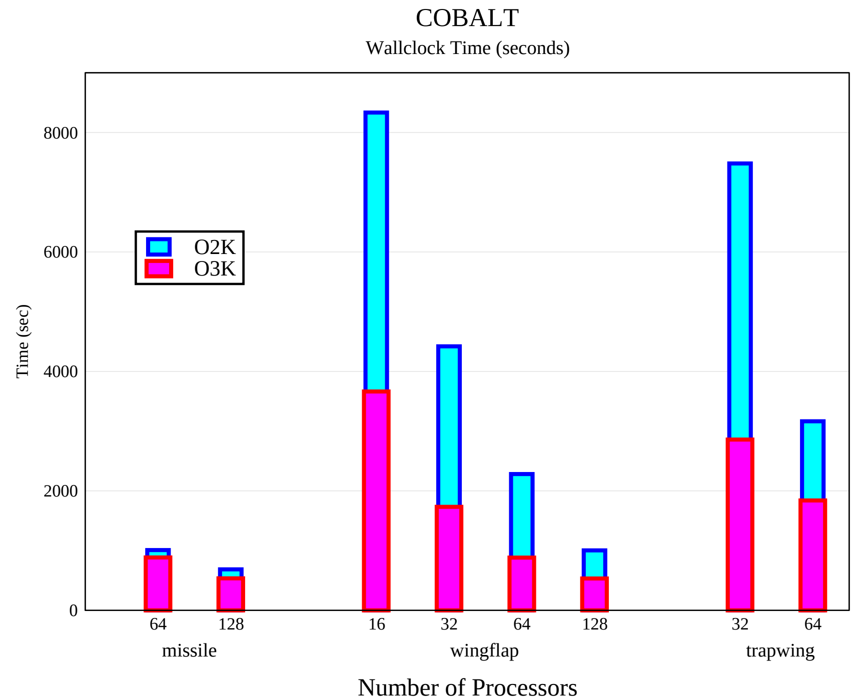

Figure 2: Timing comparisons for COBALT runs.

Comparisons of the benchmark times for the COBALT tests are given in Figure 2. On the smaller input problem "missile," the O3K does not demonstrate the expected speedup over the O2K. It is possible that for this test the initialization and overhead are large enough to prevent a significant speedup. But on the larger, more complicated runs "wingflap" and "trapwing" the O3K does significantly better.

Note, however, that since some of the COBALT tests were not performed as dedicated runs on the O2K, this may slightly exaggerate the improvement of the O3K for those tests. Indeed, this probably accounts for the much higher speedup measure of 2.46 reported for the "wingflap" input decks.

4 Conclusions

The Origin 3000 consistently outperforms the Origin 2000 with a speedup factor of slightly greater than two on a wide range of benchmarks taken from actual application codes. The enhanced performance is due to the faster clock speed and improved interprocessor communication.

It is reasonable to conclude that the improved bandwidth for communication in the O3K avoids a potential bottleneck. If the internode communication requires too much time, then the improved calculation speed due to a faster chip could be wasted because too much time would have been spent in communication. It would appear that the O3K is equipped with significantly greater bandwidth, and that this supports the faster processors.

5 Acknowledgements

This work was supported in part by a grant of computer time from the DoD High Performance Computing Modernization Program at the ERDC MSRC, Vicksburg, MS., and the ASC MSRC, Wright Patterson Air Force Base, OH.

In particular, we would like to thank Bernhard Winfried of the ASC MSRC and Marty Bullock and George Moncrief of the ERDC MSRC for their assistance in helping us set up the test problems at their sites.

Special thanks go to all of the developers of the codes and the users who contributed the input decks.

References

[1] For more information, go the the ASC website at `http://www.asc.hpc.mil/`.

[2] For more information, go the the ERDC website at `http://www.wes.hpc.mil/`.

[3] For specifications and data sheets on these computers, visit SGI's website at `http://www.sgi.com`.

[4] D. Duffy, M. Fahey, R. Fahey, J. Hensley, T. Oppe, W. Ward, and R. Alter, *TI01 Department of Defense Benchmark,* in preparation.